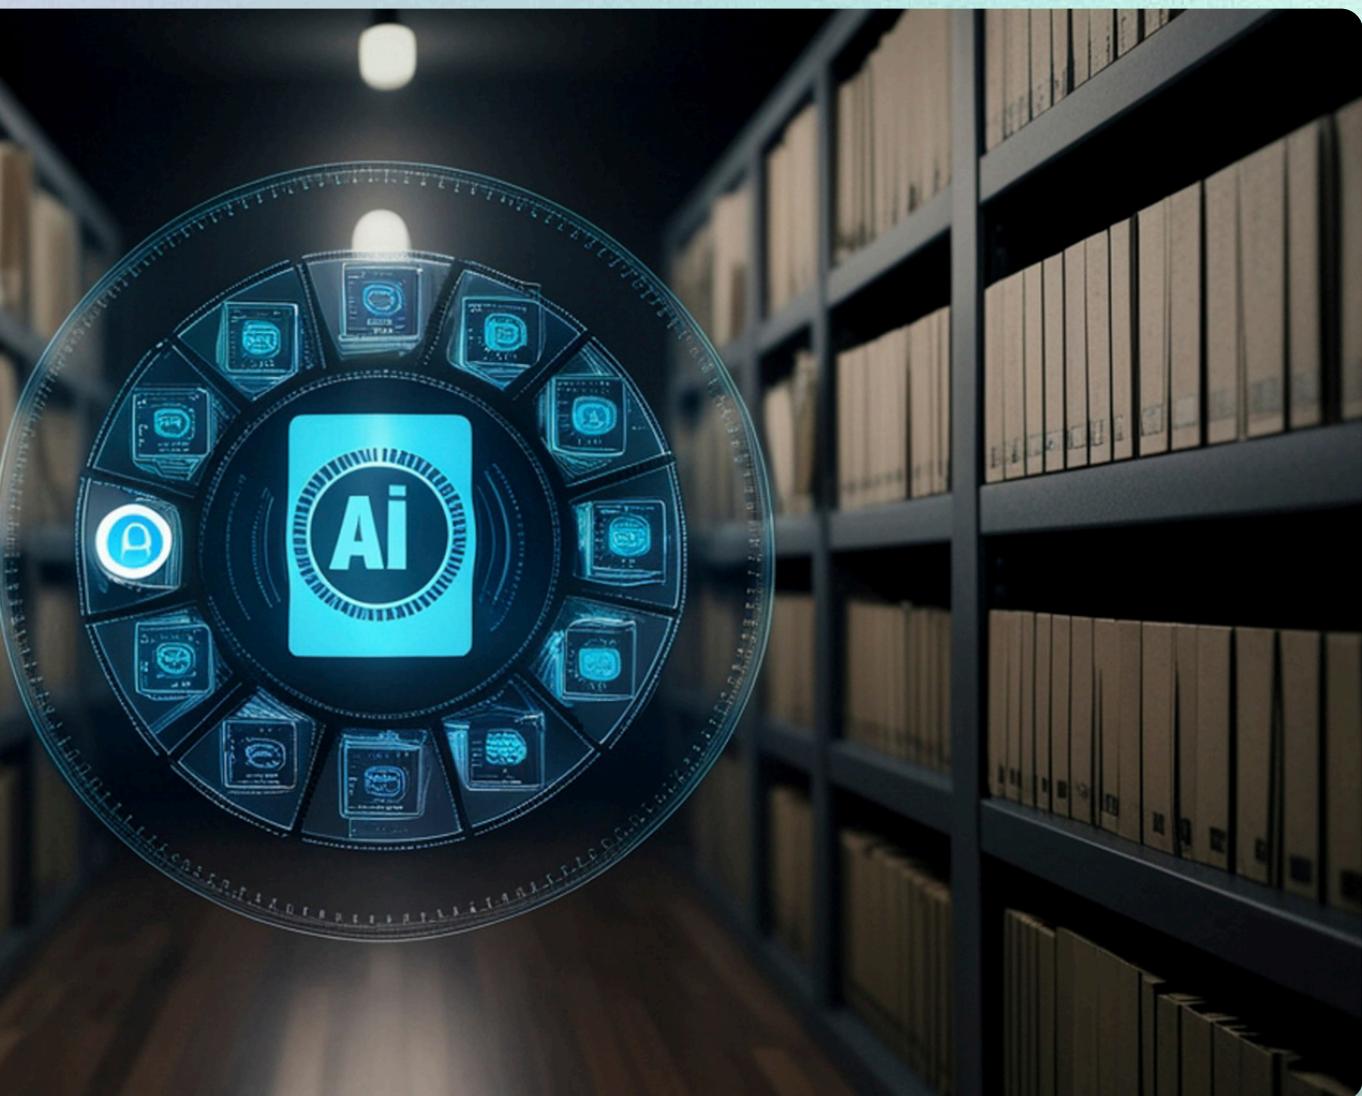


Using AI for investigations & intelligence, responsibly.

5 expert analyst tips (with actual prompts you can use)!



Introduction

We've all been there. Either overwhelmed at the beginning of an investigation, stumped in the middle of one, or too lazy to fill in the final sections of an almost-closed case. In those moments it's tempting to turn on ChatGPT and have it generate your problems away.

Is this the best way to use AI in investigations or intelligence, though? Are we avoiding risk and getting the most value out of it? For the most part, no. That's because ChatGPT, Claude and other large language models (LLMs) still aren't capable of the reasoning, association and creative thought needed to really crack investigations. If you don't believe me, just check out some fascinating research from Apple on just this topic - LLMs don't reason, but rather identify and match patterns incredibly effectively. But even with all that power, LLMs still can't replicate the kind of real logic and reasoning investigations demand. So if LLMs have such real constraints, what's the point? We believe they're still incredibly useful across investigations and intelligence. Here are 5 of the best ways to put them to work:

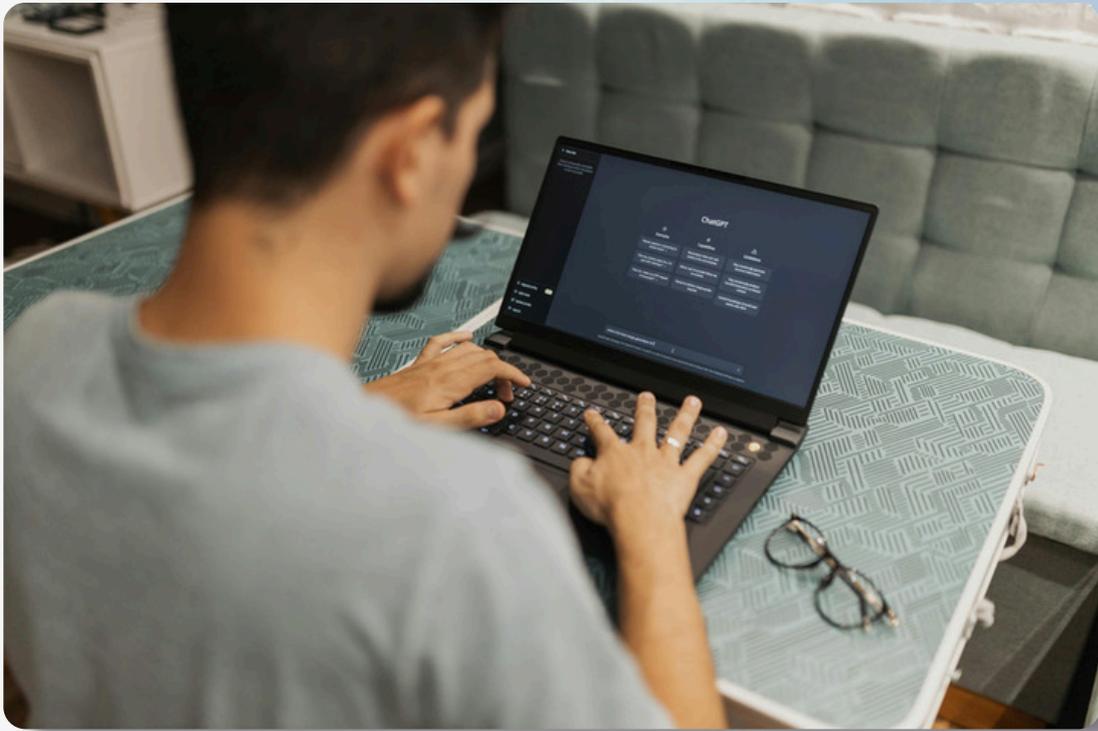
The first is the most foundational: checking our own work.

1. Red Teaming / Analytical frameworks / Structured analytic techniques

One of the most useful capabilities of modern AI models is their ability to generate alternative hypotheses at high speed. When an analyst is trying to avoid premature closure or explore a wide range of possible explanations, an AI assistant can rapidly surface multiple angles that might otherwise be overlooked.

AI is also well suited for applying structured-analytic techniques. Whether it's the Analysis of Competing Hypotheses (ACH), Key Assumptions Checks, or other systematic frameworks, AI can run through each step methodically and help the analyst organize their reasoning. This makes it easier to identify assumptions, articulate possible explanations, and outline relevant evidence.

Additionally, AI can create practical tools—such as bias-mitigation checklists—that support more rigorous thinking. These outputs can serve as quick reminders for analysts to re-evaluate their assumptions, check for cognitive traps, or review alternative explanations.



Where AI can mislead you

Despite its strengths, AI has notable limitations that can create risk if not properly managed. The most discussed issue is hallucination: AI may generate explanations or hypotheses that sound plausible but have no real evidentiary basis. Without careful review, these can slip into an analyst's reasoning. This has gotten much better as of late with recent technologies such as RAG (retrieval augmented generation) which improve LLM accuracy, but are far from perfect and still much to be double-checked.

AI also struggles to appropriately weight subtle, context-specific factors—precisely the kinds of clues that experienced analysts pick up through years of domain knowledge. It tends to generalize or smooth over unique signals that a human would immediately recognize as significant.

Best practices for using AI thoughtfully

The safest way to incorporate AI into analytical work is to treat it as an augmentation tool rather than an autonomous decision-maker. Let the model generate initial hypotheses or analytic products, but require human verification, sourcing, and refinement before accepting anything as valid. In other words, AI should spark thinking—not conclude it.

Every claim or hypothesis produced by AI should be tied back to a source. Analysts should insist on explicit citations or evidence and treat unsupported statements as provisional drafts rather than established facts.

It's also wise to run the AI through multiple structured techniques to compare outputs. For example, you might ask it to propose five potential hypotheses, then have it rank them by likelihood using a structured framework or evaluate each against available evidence. Divergent outputs can reveal blind spots or highlight which explanations withstand more scrutiny.

Finally, maintain an audit trail of prompts and responses. Keeping a record of the reasoning process is essential for transparency, peer review, and after-action analysis. It allows others to understand why certain conclusions were reached and what role the AI played in shaping the thinking.

Despite its strengths, AI has notable limitations that can create risk if not properly managed. The most discussed issue is hallucination: AI may generate explanations or hypotheses that sound plausible but have no real evidentiary basis. Without careful review, these can slip into an analyst's reasoning. This has gotten much better as of late with recent technologies such as RAG (retrieval augmented generation) which improve LLM accuracy, but are far from perfect and still much to be double-checked.

AI also struggles to appropriately weight subtle, context-specific factors—precisely the kinds of clues that experienced analysts pick up through years of domain knowledge. It tends to generalize or smooth over unique signals that a human would immediately recognize as significant.



Quick prompt template

"Using the ACH method, produce 5 mutually exclusive hypotheses to explain [incident summary]. For each hypothesis:

- 1) One-sentence description
- 2) Key evidence that would support it (list)
- 3) Key evidence that would disprove it (list)
- 4) Highest-risk assumption
- 5) Next investigative step to test it.

Do NOT invent sources; mark any items that are inference rather than sourced."

Red-team trick

- Ask the model to argue against your leading hypothesis, forcing counter-evidence and alternative logic.
- Ask the model to find logical flaws, biases or issues in sourcing.

2. Data analysis

Modern AI systems excel at providing rapid, exploratory overviews of structured datasets—especially formats like CSV or TSV files. When you're trying to get a quick sense of what's inside a dataset, AI can help outline key variables, summarize distributions, and highlight general patterns. This makes the early stages of analysis significantly faster. One of the most effective uses of AI that exists, in my past experience, is using AI to help you manage your data effectively. This means merging formats of data, normalizing data files, converting data types and even using AI to structure unstructured data (such as audio to text or tables).

AI is also effective at generating working code snippets in languages such as Python, particularly when using libraries like pandas, numpy, matplotlib, or seaborn. Whether you need to clean messy columns, perform routine aggregations, or spin up a simple visualization, the model can often produce a useful starting point. For routine tasks, this can save considerable time. New AI capabilities such as Open AI's agentic models enable models to do this for you autonomously, and new browsers and other tooling will soon enable you to run that code locally via the model without having to open an IDE.



Where AI struggles (and why it matters)

Despite these strengths, AI has notable weaknesses that require careful human oversight. For one, it may recommend statistical methods that are inappropriate for edge cases or specific data configurations. A test that works well under standard assumptions may be misleading when the underlying data violates those assumptions.

AI also tends to overlook data quality problems such as missing values, outliers, duplicates, or sampling biases. Because it generally assumes clean inputs unless told otherwise, it can generate analysis paths that appear perfectly reasonable but rest on flawed foundations.

Lastly, while effective, it can be annoying to get AI to do exactly what you want with many data file types. Don't give up though, it just takes a bit of time and effort to dial in the prompt.

Best practices for using AI in data analysis

To use AI effectively and safely, it's important to adopt guardrails that ensure the integrity of your analysis.

Start with data quality, always. Before running any AI-generated code, examine distributions, missingness patterns, duplicate rows, and outliers. This establishes a correct understanding of the dataset's limitations.

- Treat AI-generated code as a draft—not an executable truth. Ask the model to produce scripts you can run locally, then review them line by line. Ensure each transformation is intentional and that the logic matches your analytic goals.
- Demand clear statistical outputs. Whenever AI suggests an analysis, require it to specify effect sizes, confidence intervals, and—when appropriate—p-values. This discourages overemphasis on binary significance and encourages richer interpretation.
- Favor transparent models when introducing more variables. For multivariate analysis, begin with interpretable approaches such as logistic or linear regression before jumping to complex, black-box algorithms. Transparent models make it easier to audit assumptions, communicate findings, and identify errors early.

Example prompt (code generation)

```
"I have a CSV with columns: timestamp, user_id, event, lat, lon, value. Provide a Python (pandas) notebook snippet to:  
- load and clean (drop duplicates, handle missing lat/lon)  
- aggregate events per user per day  
- run a simple time-series plot of daily event counts  
Include comments and sanity-check assertions."
```

3. Visualization

AI is particularly effective at proposing suitable chart types and layouts based on the nature of the data and the intended audience. Whether you need a simple bar chart, a panel of small multiples, a network diagram, or a map, the model can quickly outline visualization options that match your communication goals, all the while saving you a ton of effort in Excel or other tools to get your data ready.

AI also excels at generating code for a wide range of visualization libraries—such as matplotlib, plotly, d3.js, or TikZ. This is especially helpful when you need a starting template, are unsure of syntax, or want to avoid manually debugging boilerplate plotting code. With a few prompts, you can get a functional script that you can refine and customize.

In addition, AI can produce clear figure captions, labels, and annotations. This includes descriptive text, callouts pointing to significant features, and narrative-style explanations that guide readers toward the intended interpretation.



Where AI can mislead you

Despite its strengths, AI sometimes recommends visual designs that are aesthetically appealing but analytically problematic. A common example is the use of truncated or irregular axes, which can exaggerate small differences or distort the apparent size of effects. Without careful scrutiny, such recommendations can unintentionally create misleading impressions.

AI may also suggest color palettes, encodings, or plot layers that obscure meaningful patterns. Overplotting, poor contrast, or too many categories mapped to the same channel can cause important differences to be lost in visual noise. In complex visualizations, these issues can reduce clarity and confuse the audience.

Lastly, it's also quite difficult on some occasions to get exactly what you're looking for with AI-generated visuals. It can be quite frustrating to edit them manually or continually reprompt, but overall - the effort is worth it.

Best practices for using AI in data visualization

The safest way to incorporate AI into analytical work is to treat it as an augmentation tool rather than an autonomous decision-maker. Let the model generate initial hypotheses or analytic products, but require human verification, sourcing, and refinement before accepting anything as valid. In other words, AI should spark thinking—not conclude it.

To ensure that AI-supported visualizations remain ethical, clear, and analytically sound, it helps to follow structured best practices.

Prioritize clarity over aesthetics. Every figure should have labeled axes, appropriate units, and clear indicators of sample size or N. These basics reduce ambiguity and enhance interpretability.

Use small multiples or faceting for comparisons. Rather than layering many series in a single plot—which often leads to clutter—break the data into side-by-side panels. This approach preserves readability and makes comparisons easier.

Check for visual integrity. Confirm that the visualization does not distort the magnitude of effects or introduce deceptive scaling. Make sure axes start at appropriate baselines and that transformations are clearly disclosed. Create alt text and an interpretive caption. Every figure should include accessible text that explains what the viewer is seeing and what the main takeaway should be. This not only improves accessibility but also reinforces analytical transparency.

Example prompt (code generation)

```
"Suggest the best 2 visualizations to compare daily event frequency across 3 regions and provide Plotly code to produce them. Include titles, axis labels, and a one-sentence caption that explains the take-away."
```

4. Deep research (final comparison)

One of AI's most useful capabilities is its ability to rapidly assemble comparative summaries from many sources. When given a set of topics, entities, or claims to analyze, AI can quickly outline how different sources treat the same issue, highlight points of convergence and divergence, and propose candidate explanations for those differences. This can dramatically shorten the time needed to produce an initial comparative framework.

In addition to comparing sources, AI is effective at generating chronological structures such as timelines. When synthesizing findings from diverse documents, reports, or datasets, AI can propose a logically ordered sequence of events and package them into concise, readable formats. In early stages of research, this helps analysts identify gaps, inconsistencies, or patterns that warrant deeper investigation.

These strengths make AI a powerful tool for generating "first drafts" of comparative or integrative analyses—documents that analysts can refine, interrogate, and substantiate through further research.



Where AI falls short — and why it matters

Despite its ability to produce polished comparative summaries, AI models are not reliable sources of factual truth. They may hallucinate citations, misattribute claims, or generate plausible-sounding but nonexistent references. These errors are often subtle and packaged in authoritative language, making them difficult to detect without deliberate verification.

In high-stakes environments—policy analysis, investigative research, scientific assessment—such misattributions can cause real harm. AI-generated summaries should never be treated as replacements for primary-source review. At best, they serve as scaffolding for human researchers; at worst, they can introduce unrecognized inaccuracies that skew downstream conclusions.

Because comparative work often depends on traceability—knowing exactly which source said what and when—any AI involvement must be matched with rigorous human checks.

Best practices for high-reliability AI-assisted research

To integrate AI safely into comparative analysis, teams should establish disciplined workflows that emphasize verification and traceability.

A strong approach begins by asking AI for a **research-first draft**—a structured attempt at comparing sources or assembling a timeline. This draft is merely a set of hypotheses about what the sources say. Analysts must then verify every factual claim directly against primary materials.

To make this process easier, AI should be required to produce inline citations and short, item-level confidence scores for each assertion. Human reviewers can then triage the claims: high-confidence items may simply require spot-checking, while low-confidence ones merit deeper inspection.

Because models sometimes repeat the same errors across prompts, it is wise to use multiple independent search tools or AI systems and cross-check their outputs. When different tools converge on the same claim—and it is confirmed in the underlying sources—the result is more trustworthy.

For high-stakes or decision-critical findings, analysts should insist on at least two corroborating, independent primary-source confirmations before treating any claim as validated. This standard mimics practices in investigative journalism, intelligence analysis, and academic research, where redundancy reduces the risk of adopting erroneous information.

1. A disciplined verification process ensures that AI accelerates research without sacrificing quality. A practical workflow might look like this:
2. AI produces a draft comparison of events, claims, or entities, along with candidate sources and confidence levels.
3. The analyst retrieves each cited source and checks the key statements, figures, or dates. Direct confirmation from primary documents is essential.
4. The draft is annotated with “verified” or “unverified” flags, along with notes identifying discrepancies or ambiguities.

This workflow preserves the speed benefits of AI while ensuring analytical integrity. It allows analysts to take advantage of AI's ability to structure information quickly, without outsourcing the fundamental responsibility of verifying truth.

Prompt

```
"Create a comparative summary of [topic] including: key facts, timelines, 3 primary sources per major claim, and a confidence level (High/Medium/Low) with justification. Highlight anything needing primary-source confirmation."
```

5. Media analysis (geolocation, image analysis, etc.)

AI can be extremely helpful in guiding the process of geolocation, even if it cannot independently confirm authenticity. When presented with an image or video, AI can outline the analytic steps an investigator should take—what to check first, what features may matter, and which tools might help. This includes identifying potential landmarks, suggesting initial hypotheses, or flagging visual cues that are worth closer examination.

One of AI's strengths lies in recommending appropriate open-source tools and methodologies. For example, it can point analysts toward tools for extracting image metadata, running EXIF inspections, conducting shadow or sun-angle analysis, or comparing topographic features. These suggestions help new practitioners learn the workflow and give experienced analysts quick reminders of techniques worth applying.

In other words, AI excels at acting as a guide—helping structure an investigation, propose an analysis plan, and highlight relevant techniques—while the human investigator performs the substantive verification.



Where AI fails—and why this matters

Despite its useful guidance, AI cannot reliably authenticate images or video. It cannot distinguish between genuine media and sophisticated synthetic or manipulated content, and it often presents tentative assessments with unwarranted confidence. This tendency to overstate certainty can mislead analysts who are not cautious.

Additionally, AI is not a replacement for specialized forensic tools. Professional EXIF forensics, photo-matching software, compression-error analysis, error level analysis (ELA), and pixel-level anomaly detection require dedicated tools and expert interpretation. AI-generated descriptions or guesses cannot substitute for these techniques.

Because media verification is often used in contexts where accuracy is paramount—journalism, human rights investigations, legal proceedings—AI-assisted analysis must always be paired with rigorous technical verification.

Best practices for responsible AI-assisted geolocation

The safest approach to using AI in media verification is to treat it as a planning assistant, not an authentication engine. Start by asking AI to outline a reproducible analysis plan describing what to check, how to measure visual indicators, and what tools to use. Then conduct the analysis manually using trusted software and primary evidence.

Preserving original files and metadata is also critical. Analysts should work from untouched media, maintain file integrity, and avoid conversions, compressions, or modifications that could destroy crucial forensic traces. Any geolocation effort should rely on multiple independent techniques. These may include cross-referencing landmarks, conducting shadow and sun-angle analysis based on known timestamps, examining EXIF metadata, comparing vegetation or architectural features, or searching for matches in online imagery archives and satellite platforms.

Equally important is maintaining a chain of custody—recording what was done, when, how, and with which tools. This documentation ensures transparency and makes it possible for other investigators or reviewers to audit the methodology.

A practical checklist for geolocation verification

A structured checklist ensures consistency and reduces the risk of overlooking key indicators. A comprehensive geolocation workflow may include:

- Preserve the original file and record a checksum. Work from a verified copy and maintain strict file integrity.
- Extract EXIF data and validate timestamps. Look for inconsistencies, missing entries, or indicators of alteration.
- Identify distinctive features such as landmarks, vegetation types, building shapes, terrain forms, or signage.
- Analyze shadows and sun position to see whether the lighting matches the claimed date, time, and geographic location.
- Search imagery archives—including street-level photos, crowd-sourced image libraries, and satellite datasets—to find potential matches or contradictions.

Following these steps does not guarantee certainty, but it substantially increases the rigor and transparency of the geolocation process.

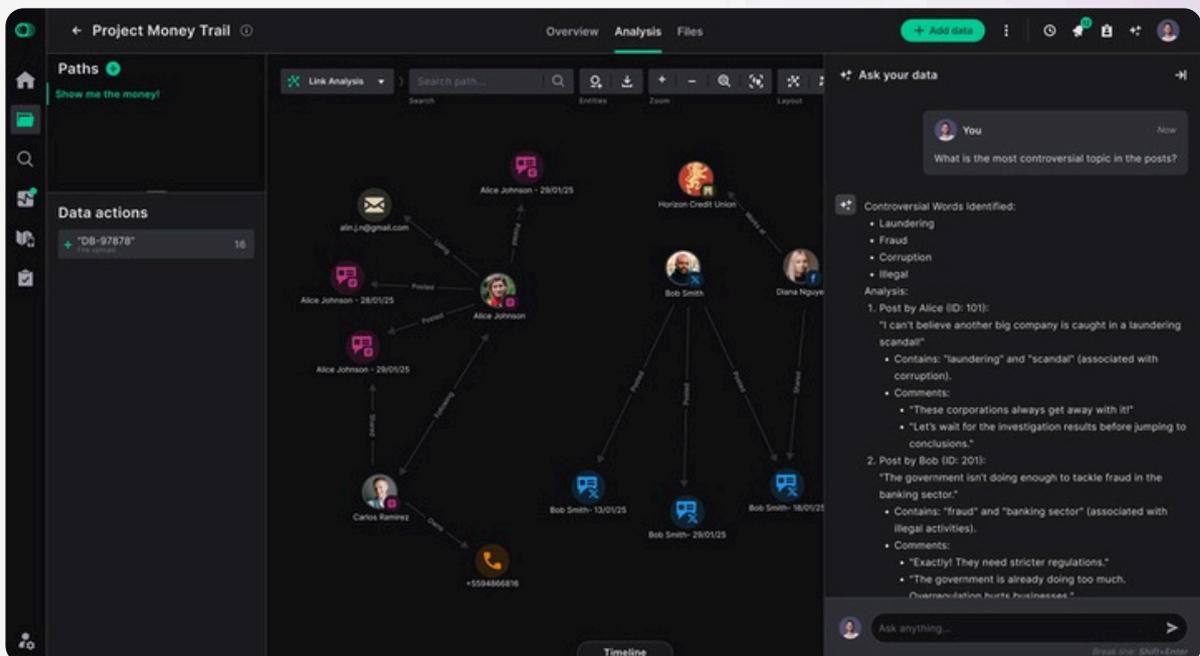
Prompt

"Given an IP-quality image and a timestamp, produce a step-by-step geolocation plan listing tools, what to extract, and how to estimate sun angle from shadows. Also list three pitfalls that could produce false matches."

Final quick checklist (apply before trusting AI output)

- Is every factual claim sourced? If not, flag as hypothesis.
- Do humans validate key steps? (Yes → proceed; No → validate.)
- Is there a reproducible audit trail (prompts + responses + run code + raw data)?
- Are alternative hypotheses considered and documented?
- For media evidence: preserve originals and use at least two independent verification methods.

One platform. Everything you need to run an investigation.



With an AI assistant that analyzes, connects, and surfaces insights like a seasoned analyst.

[Request a demo](#)